COMPAQ
White Paper

October 2001
15F2-0801A-WWEN
Prepared by:
Enterprise Storage Group
Compaq Computer Corporation

Contents

Features and Benefits of HSG80 ACS 8.6P Data Replication Manager

Abstract: The Compaq SANworks Data Replication Manager (DRM) is controller-based data replication software for disaster tolerance and data movement solutions. DRM allows all data to be mirrored between two different storage arrays, which can be in separate geographic locations. Properly configured, DRM can be a complete disaster-tolerant storage solution that guarantees data integrity in case of a storage subsystem or site failure.

Notice

Introduction

Data Replication Manager (DRM) provides remote disk mirroring by connecting primary RAID systems to remote storage RAID systems. In the event of a disaster at the primary site, a copy of the data is safe on the remote storage system and available for immediate connection to an alternate host.

Features of the DRM solution include:

- Online and real-time sequential data replication to a local or remote site
- Heterogeneous support of Compaq *OpenVMS*™, Compaq *Tru64*™, IBM AIX, Microsoft Windows 2000, Microsoft Windows NT, Novell NetWare, and SUN Solaris on same fabric or controller
- Business continuance volumes
  - Cloning at initiator and target sites
  - Snapshot support at target site
  - Remote mirrors between both sites
- Large SAN support
  - Up to 20 switches with cascaded switch supporting, for example, 8 storage arrays and 96 servers per site
  - Fibre Channel switch zoning support
- Multiple Intersite Link technologies
  - Unlimited distance with Fibre Channel-to-ATM (Asynchronous Transfer Mode) connectivity with line speeds of T1/E1 through OC3 using the Open Systems Gateway (OSG)
  - Unlimited distance with Fibre Channel-to-IP connectivity by means of 10/100 Ethernet with Gigabit Ethernet expected by the end of 2001
  - Up to 100 km (62 miles) with Wave Division Multiplexing (WDM)
  - Up to 100 km with Very Long Distance Gigabit Interface Converters (GBICs) or up to 10 km support with longwave GBICs
- Support for both synchronous and asynchronous replication modes
- Write history logging and mini-merge reconstruction
- Cluster and stretched cluster high-availability solutions
- Non disaster tolerant LUN support at initiator and target sites
- Active/active bi-directional solutions
- Dual and Single Switch and Single Fabric entry level configurations for development and testing
- Up to 12 remote storage sets per pair of arrays
- Up to 12 local (non remote copy) storage sets per array

- Up to 12 remote copy sets per association set
- New and improved procedures for AIX, OpenVMS, NetWare, Solaris, Tru64, Windows 2000, and Windows NT
  - Ten documented failover and failback procedures covering planned and unplanned contingencies.
  - Documented procedures for moving data, performing multisite replication, and growing a storage set.
  - All procedures are supported by the full DRM dual fabric, four-switch minimum solution and the limited two-switch, dual fabric DRM Dual Switch/Single Site solution.
- The following application notes and white papers are available at

  www.compaq.com/products/sanworks/drm/documentation.html
  - *Maximizing Efficiency Through Data Replication Manager Data Movement*
  - *DRM over an Internet Protocol Link*
  - *Data Replication Manager Solutions Design Guide*
  - *DRM and Virtual Replicator for Windows 2000/NT Application Note*
  - *Using Data Replication Manager with Oracle8i Under Tru64*
  - *Oracle Storage Compatibility Testing-Remote Mirroring Using Compaq SANworks Data Replication Manager*
  - *Data Replication Manager with VERITAS Volume Manager 3.x and Sun Solaris Application Note*

The benefits of the DRM solution are discussed with the description of each feature in the "DRM Features and Benefits" section of this paper. To completely benefit from the DRM solution, go to the "Supported Hardware and Software" section to verify the necessary hardware and software.

DRM has tested and is qualified with:

- *SANworks*™ Command Scripter 1.0
- SANworks Secure Path (various versions by operating system)
- Large LUN support for Windows 2000 and Windows NT 4.0
- *StorageWorks*™ Virtual Replicator
- SANworks Enterprise Volume Manager
- Veritas Cluster Services and Veritas Volume Manager for Solaris

DRM 8.5P is supporting and 8.6-1P will soon test with:

- Oracle 8*i,* by means of the Oracle Storage Compatibility Program in synchronous and asynchronous modes
- SAP
- *ProLiant*™ HAF500 clusters and related high-availability services

SANworks Data Replication Manager

The Compaq SANworks HSG80 ACS 8.6P Data Replication Manager is data replication software for disaster tolerance and data movement solutions. The software works with the StorageWorks Fibre Channel MA8000/EMA12000/EMA16000 and RA8000/ESA12000 storage systems.

DRM allows all data to be mirrored between storage elements in two different storage arrays that can be in separate geographic locations. Each write I/O is sent to both sets of storage; reads occur only at the local storage. Properly configured, DRM can be a complete disaster-tolerant storage solution that guarantees data integrity in case of a storage subsystem or site failure.

DRM is a solution for copying data online and in real time to remote locations through a local or extended Storage Area Network (SAN).

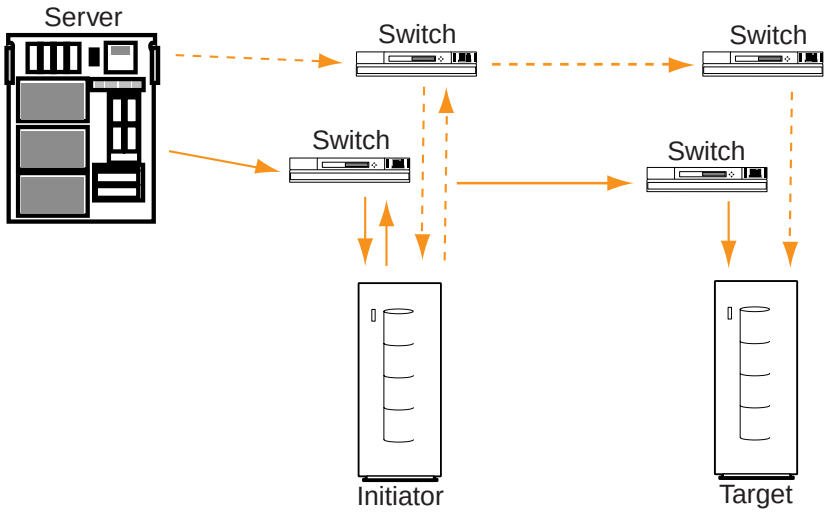

Figure 1: SANworks Data Replication Manager overview

For more information on DRM, review the additional documents available on the web, such as:

- *Data Replication Manager HSG80 ACS Version 8.6-1P Configuration Guide*
- *Data Replication Manager HSG80 ACS Version 8.6-1P Failover/Failback Procedures Guide*
- *Data Replication Manager HSG80 ACS Version 8.6-1P Scripting User Guide*
- *Data Replication Manager HSG80 ACS Version 8.6-1P Application Notes*
- *Disaster Tolerant SAN Solution Description*

All of the DRM documents, including this paper, are updated frequently. If this paper or other DRM documentation is older than four months, updated documents may be available at

www.compaq.com/products/sanworks/drm/documentation.html

DRM Features and Benefits

This section describes the features of DRM and the benefits they provide to the DRM solution.

Disaster tolerance—The redundancy of the DRM hardware provides disaster tolerance. Each Modular Array (MA) has dual RAID controllers, mirrored cache, dual cache batteries, and dual power supplies. SANworks Secure Path software or native multipath driver support from the operating system deliver a multipath driver that automatically fails over to an alternate data path when dual intersite links and controllers are used in the network.

Disaster recovery—In the event of a disaster, the Compaq SANworks DRM ensures that an online copy of data is available at the alternate site to support rapid resumption of critical processing at the alternate site. Resumption can usually occur within minutes, instead of the hours or days that other disaster recovery plans can take. Failover and failback processing may be done manually or by an automated script to reduce errors.

Moving data with DRM—The Compaq SANworks DRM can replicate, by remote mirroring, large quantities of data for data warehousing and data mining applications. DRM is also suited for moving important data off of storage subsystems temporarily for system maintenance.

Heterogeneous operating systems—DRM supports compatible operating systems sharing the same controller. To be compatible, all of the operating systems must support the same level of SCSI command and control. The following table lists the supported operating systems and the SCSI command level that each operating system supports. Additional information on the sharing of Compaq storage subsystems is available in the *Heterogeneous OpenSAN Design Guide* available at

www.compaq.com/storage/san_index.html

Table 1. Heterogeneous Operating Systems

Operating System	Versions	SCSI-2	SCSI-3
Compaq OpenVMS	7.2-1H1, 7.2-2, 7.3	No	Yes
Compaq Tru64 UNIX	5.1, 5.1a	Yes	Yes
IBM AIX	4.3.3	Yes	No
Microsoft Windows NT	4.0	Yes	Yes
Microsoft Windows 2000	Svr., Ad. Svr., DC	Yes	Yes
Novell NetWare	5.1	Yes	Yes
Sun Solaris	7, 8	Yes	No

Storage-based—Data replication is performed at the storage subsystem controller level and is totally transparent to the host, alleviating unnecessary host cycles to perform the data mirroring function.

Fibre Channel speed—Data replication is performed up to full Fibre Channel speeds (100 MB/sec) across the SAN. The dual-controller design provides extra processing power to service host I/O requests. A read-ahead cache mechanism can increase performance 30 percent to 40 percent for applications that sequentially read data.

Zoning—Zoning allows a switched fabric to be partitioned into multiple zones. It limits a controller to a specified zone or zones and limits the number of controllers a server can see. This feature increases the number of servers allowed per SAN and enables the administrator to dedicate a zone for a specific function such as backup. This feature also allows the consolidation of local area SANs into one larger wide area SAN and, therefore, sharing of the storage infrastructure.

Clone and snapshot—SANworks DRM supports cloning of RAID 0, RAID 1, and RAID 0+1 storagesets at either the initiator or target site. A clone is a physical point-in-time copy of the data. DRM also supports the creation of snapshot copies of any storageset at the target site. A snapshot is a virtual point-in-time copy of the data. Clone and snapshot are functions that minimize the downtime required for system backups and data migration activities. These capabilities ensure business continuance by allowing parallel processing and are designed for customers who cannot disrupt their computing operations for management activities.

Non remote copy set—Logical storage units that are created on a storage array and are not part of a remote copy set are referred to as a *non remote copy set storage*. Remote copy sets and non remote copy sets are logical storage units that can exist on the same subsystem at both the initiator and target sites. This provides added LUN support and configuration flexibility. Non remote copy sets may be different at the initiator and target sites. Also, the creation of non remote copy sets maximizes the use of existing infrastructure and supports snapshot or cloning operations.

Bootless failover of data (noncluster quorum, nonsystem disks) for Tru64, OpenVMS, and Windows 2000—At failover initiation, the bootless failover and failback procedures allow the surviving server to dynamically recognize the new LUNs (failed over remote copy sets, snapshots, or clones) without rebooting the server.

Wavelength Division Multiplexing and Dense Wavelength Division Multiplexing technology—DRM supports WDM and Dense Wavelength Division Multiplexing (DWDM) optical networks. These network connections provide multiple high bandwidth connectivity of enterprise-level SANs over a metropolitan area network (MAN) through private or public fiber optic networks. The use of WDM-based technology allows the intersite infrastructure to expand without investment in new cable. For example, the existing fiber may support only a single intersite link, but adding WDM at each site allows the same cable to be used for dual Fibre Channel links in addition to the WAN connection between those same sites.

High availability with no single point of failure—DRM configurations ensure highly available data to host applications by having no single point of failure (NSPOF) in their hardware and software components. Dual host bus adapters (HBAs), Fibre Channel links, switches, and storage subsystems provide redundancy in hardware components. Secure Path software or similar multipath driver support in the operating system provides HBA-failover support from server to storage, which provides system redundancy from a software perspective. SANworks Secure Path supports AIX, NetWare, Solaris, Windows 2000, and Windows NT operating systems. Secure Path is not needed for OpenVMS and Tru64 because multipath failover is supplied with these operating systems.

Active/active bi-directional—The bi-directional DRM solution addresses the growing need to ensure continuous availability of applications that are critical to daily business operations. DRM is the first storage solution to enable two sites connected by switched Fibre Channel to use each other to maintain synchronized copies of online data, maximizing resource utilization while enabling business continuance, even in the event of disaster.

Windows 2000 and Windows NT Large LUN Support—After enabling the Windows Large LUN feature within the operating system, DRM supports up to 64 LUNs in SCSI-2 mode or 63 LUNs in SCSI-3 mode.

Note: A single DRM instance is limited to 12 remote copy sets and 12 non disaster-tolerant LUNs per site for a total of 36 LUNs. Even this exceeds the drive letter limit of Windows 2000 or Windows NT.

DRM Dual Switch/Single Site solution—The DRM Dual Switch/Single Site solution allows customers to implement a host-independent remote copy solution at the departmental level within the enterprise. This solution does not provide protection against most natural disasters due to its limited separation distance. This solution is supported by all DRM supported operating systems. A specific solution for the Microsoft Windows 2000 and Windows NT environment is available using the Departmental DataSafe solution available at

www.compaq.com/products/storageworks/solutions/essa/deptdatasafe.html

Extended SANs over direct Fibre Channel—DRM replicates data over direct Fibre Channel networks of up to 100 km (62 miles) by way of the Very Long Distance GBIC. Data replication can be performed at full 100 MB/sec Fibre Channel speeds to approximately 50 km (31 miles). An intersite link longer than 50 km (31 miles) cannot maintain the full 100 MB/sec data rate because of the limited number of buffer-to-buffer credits available in the current version of switch firmware. The longwave GBIC for full Fibre Channel bandwidth data transmission up to 10 km (6.2 miles) is also available.

Extended SANs over IP capability—When available, DRM can use new Fibre Channel (FC) to IP network gateways to extend the SAN across large distances using the existing IP network infrastructure. These industry-standard gateways provide the needed network interface and use IP-based tunneling technology to provide the dedicated bandwidth and data protection.

For more information on Extended SANs over IP capability, refer to the *Compaq SANworks Data Replication Manager over an Internet Protocol Link Application Note* at

www.compaq.com/products/sanworks/drm/documentation.html

Extended SANs over ATM—Disaster recovery sites can be located hundreds, even thousands, of miles from their corresponding production sites with the use of FC-to-ATM Open Systems Gateways (OSGs). The OSG provides FC-to-ATM connectivity in telecommunication networks for Compaq SANworks extended SAN-WAN-SAN solutions. ATM connectivity is designed for contingency sites that must be very long distances apart to protect against wide area disasters or where direct fiber optic connectivity is unavailable.

The FC-to-ATM gateways provide a wide range of WAN support and flexibility to meet customer requirements in the most cost-effective manner. The FC-to-ATM gateway offers an OC3 multimode fiber (MMF) physical interface and supports access speeds ranging from 1.544 Mb/sec to 155 Mb/sec. Data rates less than OC3 (155 Mb/sec) require the ATM switch (available from a WAN provider).

- **Multiple T1 lines**—With the optional ATM switch, the FC-to-ATM OSG can support one or more T1/E1 lines. Multiple T1/E1 lines provide increased bandwidth for customers who need faster data rates. Load balancing across these multiple lines is a function of the switch and the interconnect to the multiplexer.
- **Fractional ATM lines**—The FC-to-ATM OSG can support fractional lines, such as fractional T3/E3 or fractional OC3 lines. By supporting fractional ATM lines, customers can lease the exact bandwidth or access speeds necessary to support their requirements.

- **Shared ATM lines**—The FC-to-ATM OSG can also support shared or multi-use ATM lines. A shared ATM line allows only one telecommunication link between sites, but the line bandwidth is divided between the two Fibre Channel fabrics. Another common application could be DRM and non-DRM use of the same shared ATM line. A shared ATM line may be used when the customer's bandwidth requirements do not justify two separate dedicated ATM lines.
- **Data encryption**—Data can be encrypted for transmission over public ATM lines, providing companies such as financial institutions the additional data security they need.

For more information on the DRM over FC-to-ATM OSG, refer to the *Compaq Data Replication Manager over an ATM Link Application Note* at

www.compaq.com/products/sanworks/drm/documentation.html

Data Replication Process

Data from the host server is written to a storage system at the initiator site. The array controllers at the initiator site write the data locally and send a copy of the I/O request to the storage subsystem located at the target site.

The target storage system then writes the data to its local drives. As a result, the data is mirrored across two remote locations. Should the host server or storage resources at the initiator site become unavailable, processing resumes using the data copied to the target site. While active, the target of the replication is unavailable except as the source for a business continuance volume snapshot or clone.

Replication Modes

The two modes of data replication between the initiator and target sites are synchronous and asynchronous. DRM supports both modes and Compaq recommends both for performance and data integrity.

Synchronous Transfer Mode

Synchronous mode offers the highest level of data protection because the local and remote copies are identical at all times. Synchronous mirroring ensures that data copies are always identical, preventing critical data loss in the event of a failure or disaster.

In synchronous mode, data is written almost simultaneously to the cache of the local and remote storage systems. The writing process occurs before the application I/O is completed, ensuring the highest degree of data consistency. As shown in Figure 2, an application running on System A will receive the I/O complete response from Data System A only when both Data System A and its mirror have the data to be written to disk.

For applications that perform I/O in a serial fashion, this means that the next write does not start until the current write is completed. There is a direct correlation between distance and I/O completion time: the longer the intersite distance the longer it takes for a replication I/O to complete.

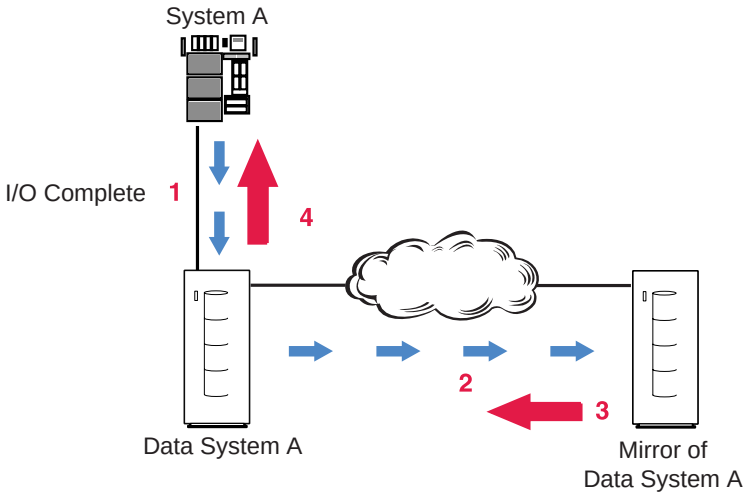

Figure 2: Synchronous transfer mode

Asynchronous Transfer Mode

With asynchronous mode, local and remote copy sets are not always identical. There is a delay between the time the local storage is updated and the time that the remote storage is updated. Asynchronous mode replication provides better response times in application environments such as those with long distances and low I/O rates. In asynchronous mode, data is written to the source system, which acknowledges and completes the I/O before synchronizing the data with the target system.

The number of outstanding unsynchronized write operations at any time is predetermined by the customer, depending on their level of tolerance for lost data. A maximum of 240 outstanding I/Os per controller is supported by the current generation of controllers. These 240 outstanding I/Os are shared among all remotely mirrored logical storage units using the pair of controllers.

Except for limited cases, Compaq does not recommend the use of the asynchronous mode because there is a risk of data loss if the local site fails before the remote site is updated. Asynchronous mode may run significantly slower than synchronous mode because the design of DRM was based on maintaining the highest level of data integrity possible.

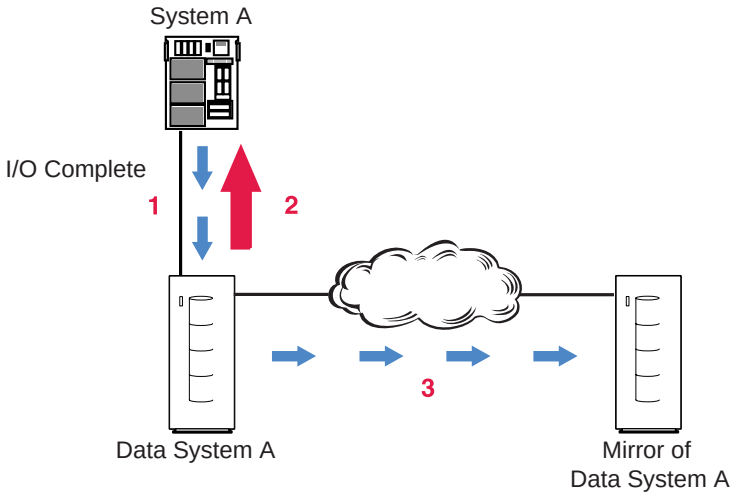

Figure 3: Asynchronous transfer mode

Remote Copy Set

A remote copy set is a bound set of two units used for long distance mirroring. Each unit can be a single JBOD disk, RAID-0 stripeset, RAID 1 mirrorset, or RAID 0+1 Stripe-Mirrorset or RAID 3/5 RAID set. In Figure 4, the LUN Remote Copy Set (RCS1) consists of Unit A and Unit C. RCS2 consists of Unit B and Unit D.

During replication, the local side of the remote copy set can be made visible to any of the hosts on the SAN. The target side of the remote copy set is owned by the local or initiator controller and is unavailable for use by a server. During failover conditions, the target temporarily becomes the initiator and the initiator becomes the target with replication going in the reverse direction. Again, only the initiator (or target running in initiator mode) is made visible to the SAN. When available, business continuance volumes, such as clones or snapshots of the target, are also visible to the SAN. These target-based snapshots and clones are useful for performing remote backups of the active databases and use non-RCS LUNs.

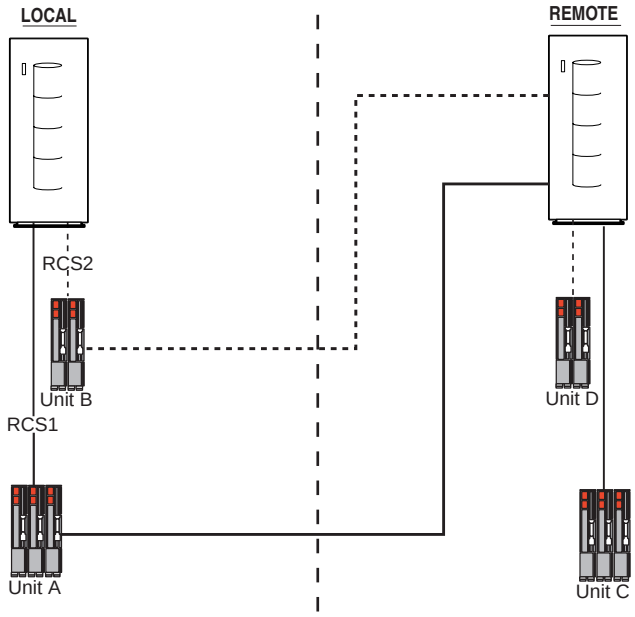

Figure 4: Remote copy sets and association sets

Association Set

An association set is the grouping of remote copy sets that can transition to the same state at the same time. The host uses association sets to keep multiple units consistent with one another. For example, when one association set member goes failsafe locked inoperative (is no longer accessible), all association set members go failsafe locked.

With write history logging mini-merge (merge of data in the write history log), write ordering is maintained for all members within an association set if the option is enabled. As shown in Figure 4, both remote copy sets are included in the one association set and will share the same write history log disk.

An association set can have up to 12 remote copy set members. All members of an association set must be on the same fabric to maintain controller cache coherency. When members are added to an association set, they are moved to the same initiator controller and fail over together. A maximum of 12 association sets, one per RCS, are available per pair of storage arrays running in DRM mode.

Basic DRM Configuration

The cabling in this configuration supports two redundant fabrics. As shown in Figure 5, the first HBA in Host A is connected to Switch C, which belongs to Fabric A. The second HBA in Host A is connected to Switch D, which belongs to Fabric B.

This dual-fabric SAN provides NSPOF at the fabric level. For example, broken cables, switch updates, or an error in switch zoning can cause one fabric to fail, leaving the other to temporarily carry the entire workload.

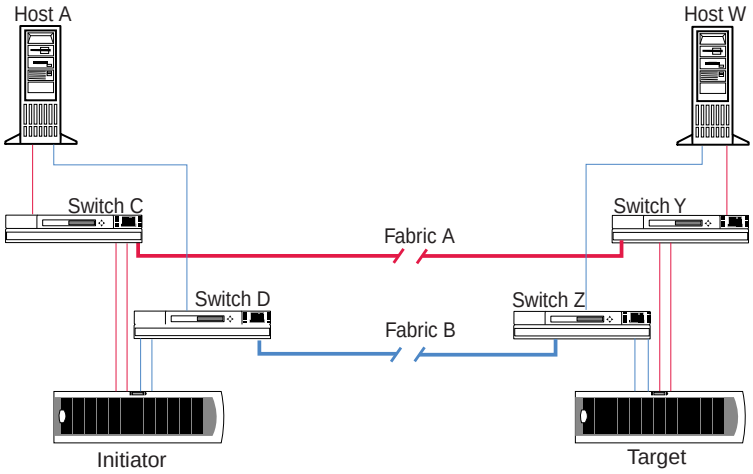

Figure 5: Basic DRM configuration

The allowable distance between the initiator and target depends on the cabling and GBIC used for the intersite links. A shortwave GBIC supports distances up to 500 meters with 50-micron multimode cable and 200 meters with 62.5-micron multimode cable. A longwave GBIC with single mode fiber supports distances up to 10 km (6.2 miles). A Very Long Distance GBIC or WDM supports intersite links up to 100 km (62 miles). As shown in the next two examples, DRM over ATM or DRM over IP is available for even longer distances or for those environments where fiber optic cable is available between two sites.

Long Distance Solutions

The following configurations explain how to use DRM over long distances.

Extended DRM over ATM Configuration

The extended DRM over ATM configuration is similar to the simple DRM configuration except for the use of Fibre Channel-to-ATM gateways from CNT Corporation. In this example a second server has been added at each site showing support for local clusters or different but compatible (at the storage level) operating systems. As shown in Figure 6, two FC-to-ATM gateways are required at each site, one for each fabric.

Due to limitations in the OSG, a DRM over ATM solution is limited to what can be connected to the available 14 ports on each of the four 16-port switches. The following options are based on one port per HBA pair and two ports per HSG80 pair per switch.

- Six storage arrays and 2 servers
- Five storage arrays and 4 servers
- Four storage arrays and 6 servers
- Three storage arrays and 8 servers
- Two storage arrays and 10 servers
- One storage array and 12 servers

As new FC-to-ATM gateways become available, support for larger DRM over ATM solutions should also become available.

DRM over ATM and the FC-to-ATM gateway support intersite link speeds ranging from 1.54 Mb/sec to 155 Mb/sec. The default connection for the gateway is 155 Mb/sec with multimode fiber. The use of another intersite connection type or speed requires the use of an ATM switch that interfaces the OC3 multimode output of the gateway into whatever is used for the intersite link. For example, if a pair of classic T1 or E1 intersite links are used, then the ATM switch must support the physical interfaces and data rates of both OC3 and T1 or E1.

DRM over ATM also supports the use of external ATM data encryption devices allowing use of DRM over public ATM networks without concern that the data may be compromised. For more information on the DRM over FC-to-ATM configuration, refer to the *Compaq Data Replication Manager over an ATM Link Application Note* at

www.compaq.com/products/sanworks/drm/documentation.html

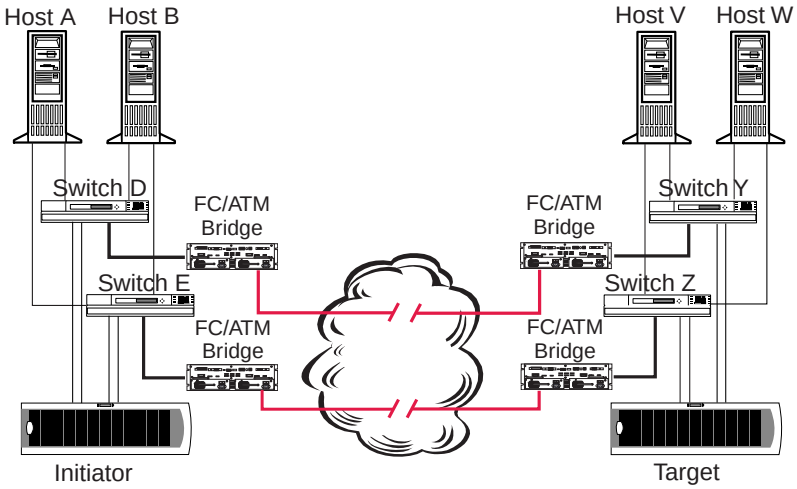

Figure 6: Extended DRM over ATM configuration

Extended DRM over IP Configuration

The extended DRM over IP configuration is similar to the simple DRM configuration except for the use of FC-to-IP gateways. As shown in Figure 7, two FC-to-IP gateways are required at each site, one for each fabric.

The FC-to-IP gateways do not have the same design limitations as FC-to-ATM gateways. The DRM-to-IP solution is limited to the same 20 switch fabrics as DRM over direct fiber solutions, which means that DRM over IP is currently limited to supporting up to 96 servers and 8 storage arrays.

Initial versions of the FC-to-IP gateway support direct connection to at least one of the current Ethernet protocols (10 Mbps, 100 Mbps, or 1 Gbps). Future versions are expected to support higher rates as they become available. The FC-to-IP gateway will use whatever intersite network bandwidth is set aside for the storage interconnect.

The IP tunnel that is created to support the FC-to-IP traffic also provides enhanced security of the data because of the nature of IP tunnels. Compaq recommends designing the IP tunnels to carry all of the intersite data traffic in case either—but not both—links fail.

For more information on the Extended DRM over IP configuration, refer to *Compaq SANworks Data Replication Manager over an Internet Protocol Link Application Note* at

www.compaq.com/products/sanworks/drm/documentation.html

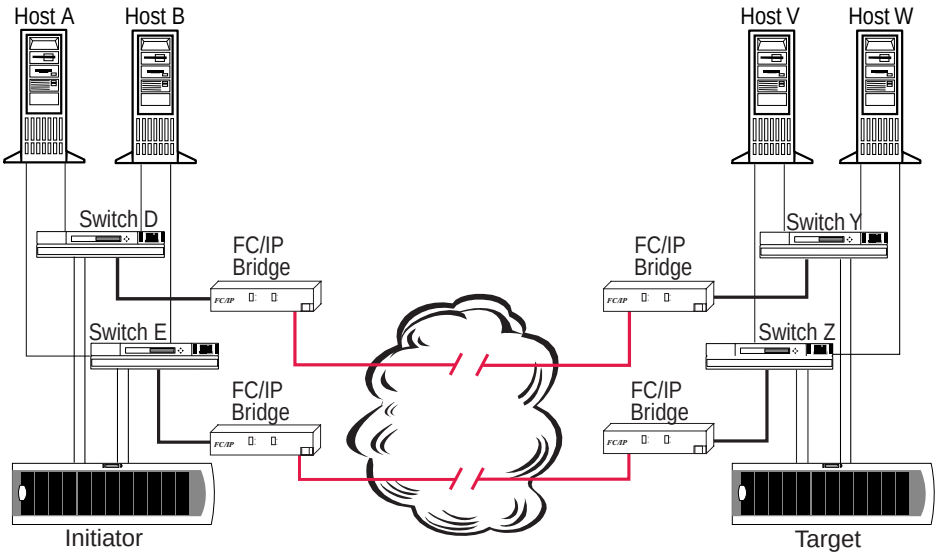

Figure 7: Extended DRM over IP configuration

Special Purpose DRM Solutions

The following sections provide different DRM configurations for special circumstances.

Bi-directional DRM Configuration

DRM supports active/active bi-directional solutions by using two sets of storage arrays, one set for each direction. This allows both sites to be actively processing data and backing up each other in the event one of the two sites fails. The Compaq Storage Solutions team created the bi-directional DRM solution, which is a turnkey deployment of the basic bi-directional DRM for Windows NT and Windows 2000. For more information on this solution, go to

www.compaq.com/products/storageworks/solutions/bd-drm/index.html

As shown in Figure 8, there are two servers at each site. These site-specific pairs of servers could be clustered during normal operations to provide highly available applications. When the failure of one of the sites occurs, the application that was running at the failed site is moved to the backup member of the surviving cluster and started using the surviving storage and the other half of the surviving cluster.

With additional servers and storage, the bi-directional DRM can scale to 20 switches in each fabric supporting, for example, 96 servers and 8 arrays at each site. All intersite technologies such as DRM over ATM, direct fiber, IP, and WDM support bi-directional use of DRM.

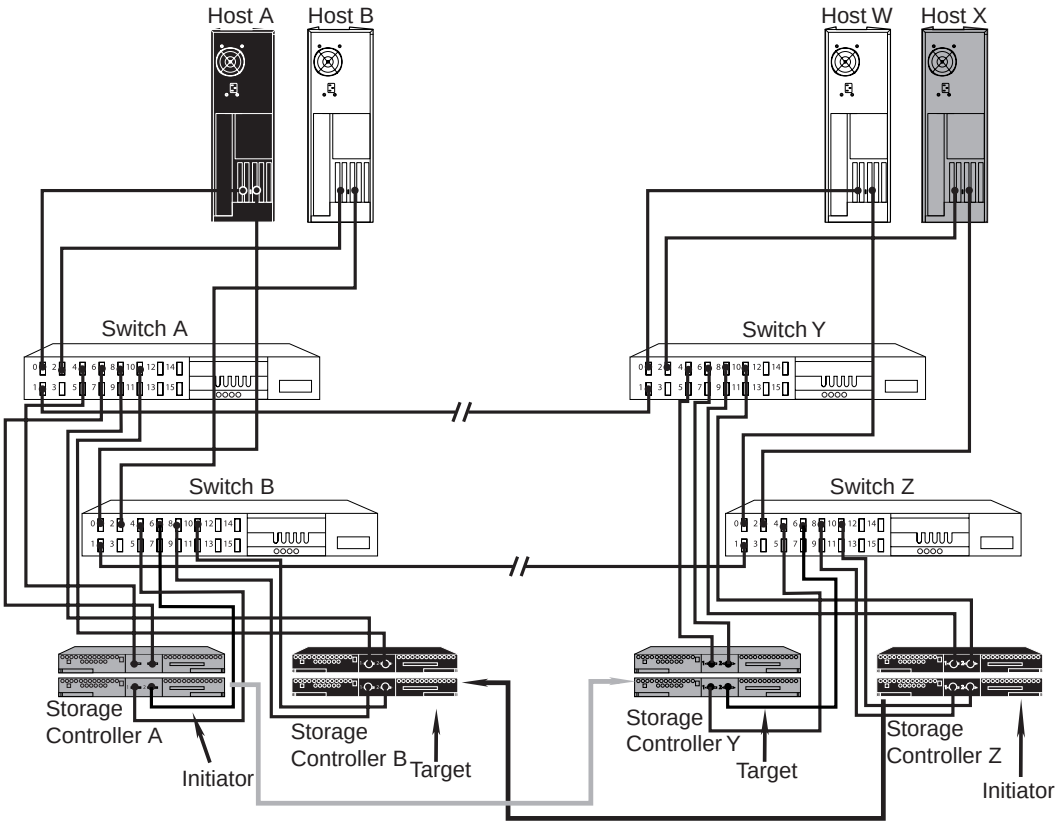

Figure 8: Bi-directional DRM configuration

Stretch Cluster Support

DRM supports stretched Microsoft Cluster Servers (MSCSs) running Windows 2000 or Windows NT. This means that half the cluster is at the primary site and the other half is at the alternate site. Should the primary server fail, MSCS fails over the application to the surviving server at the alternate site and resumes operations using the primary site storage.

Applications running in a stretched cluster in failover mode incur a performance penalty because of the time it takes to read or write data across the intersite link. This performance penalty is directly proportional to the distance between the two sites. The greater the distance between the sites, the greater the penalty. Tests have shown that almost no impact was observed with separation distances up to 100 km (62 miles).

For more information on stretch cluster support, see the Compaq ProLiant HA/F500 website at

http://www.compaq.com/solutions/enterprise/ha-f500.html

DRM Dual Switch/Single Site Configuration

Compaq supports the DRM Dual Switch/Single Site configuration for those environments that only need local data protection in the event of local disasters or that are used as local test beds for operational DRM solutions. This solution uses only two switches, where each switch creates a fabric "in a box" instead of the multiswitch fabrics supported in the non-Dual Switch/Single Site DRM solutions.

The Compaq Storage Solutions team has created the Departmental DataSafe solution for those customers requiring a turnkey data protection solution that is based on Windows 2000 or Windows NT. More information on the Departmental DataSafe solution is available at

www.compaq.com/products/storageworks/solutions/essa/deptdatasafe.html

Because of its design, a DRM Dual Switch/Single Site is limited to supporting one "switch worth" of connections (8 or 16 ports). Therefore, given a DRM Dual Switch/Single Site consisting of two 16-port switches, the maximum configuration is one of the following:

- Twelve servers and 1 pair of arrays
- Eight servers and 2 pairs of arrays
- Four servers and 3 pairs of arrays

The maximum DRM Dual Switch/Single Site configuration is slightly larger than the DRM-to-ATM configuration because there is no intersite link filling the sixteenth port like there is in the DRM over ATM solution.

The HSG80 and the server HBA use only shortwave GBICs. This means that the DRM Dual Switch/Single Site configuration is limited to 500 meters of 50 micron or 200 meters of 62.5 micron multi-mode fiber optic cable between the HBA and either switch and between the controller and either switch. This limits the DRM Dual Switch/Single Site to a maximum separation of 1 km using 50 micron cable (400 meters with 62.5 micron cable) between primary and alternate servers and primary and alternate storage arrays.

To achieve this maximum distance, Compaq recommends locating the two switches somewhere between the primary site and the alternate site. Both switches should be installed in separate locations with unique fiber paths between the switch and both sites. For example, a single backhoe will not cut both fabrics isolating the primary and alternate sites from each other. The DRM Dual Switch/Single Site is not designed to survive large area natural disasters such as tornadoes, earthquakes, or hurricanes because of the limited intersite distance. The DRM Dual Switch/Single Site may not survive or continue to operate after two unrelated failures due to its limited NSPOF configuration.

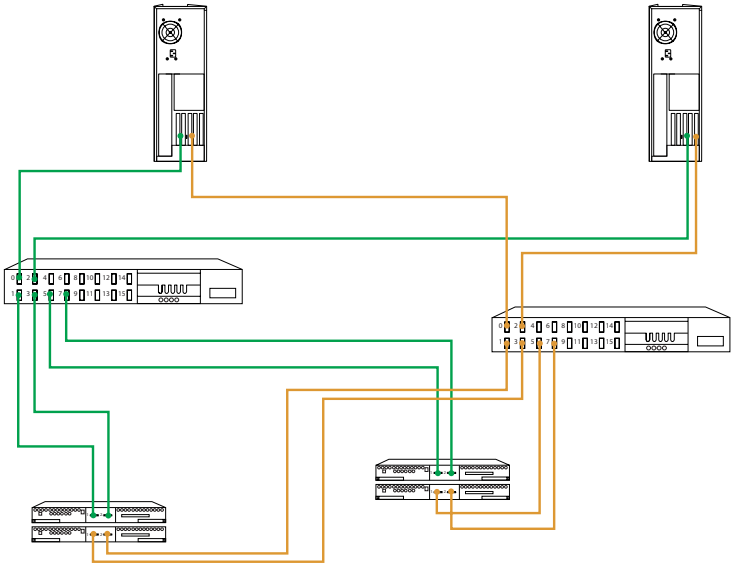

Figure 9: DRM Dual Switch/Single Site configuration

Single Switch

The Single Switch DRM solution is designed for small, single-site entry level tests and proof of concept demonstrations. This nondisaster-tolerant solution can also be used for producing copies of data needed for data migration or data mining. The maximum configuration at the time of this note is one 16-port switch, where 4 ports are used for each fabric at each site. These four ports can support a maximum of two servers and one storage array per simulated site. Fabric zoning is used to create the two logical fabrics used by DRM. An example of the solutions is shown in Figure 10.

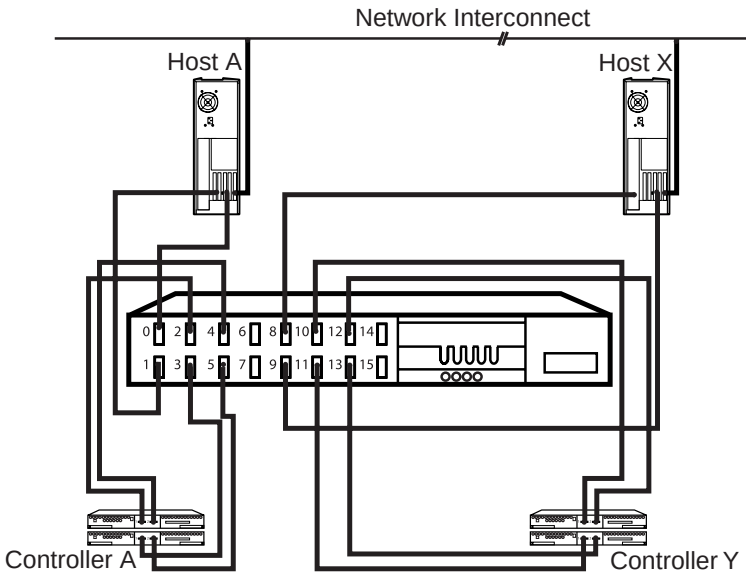

Figure 10: Single Switch

Single Fabric

The Single Fabric DRM solution is designed for small, entry level tests and proof of concept demonstrations where some distance is needed between each of the two switches in the solution. This nondisaster-tolerant solution can also be used for producing copies of data needed for data migration or data mining. Fabric zoning is used to create two logical fabrics out of the one physical fabric.

The maximum configuration at the time of this note is two 16-port switches, with one switch at each site. The two switches share one intersite link leaving up to 15 ports for servers and storage controllers (including one spare). The 14 remaining ports will support up to 10 servers and 1 array, 6 servers and 2 arrays, or 2 servers and 3 arrays per site. An example of the solutions can be found in Figure 11.

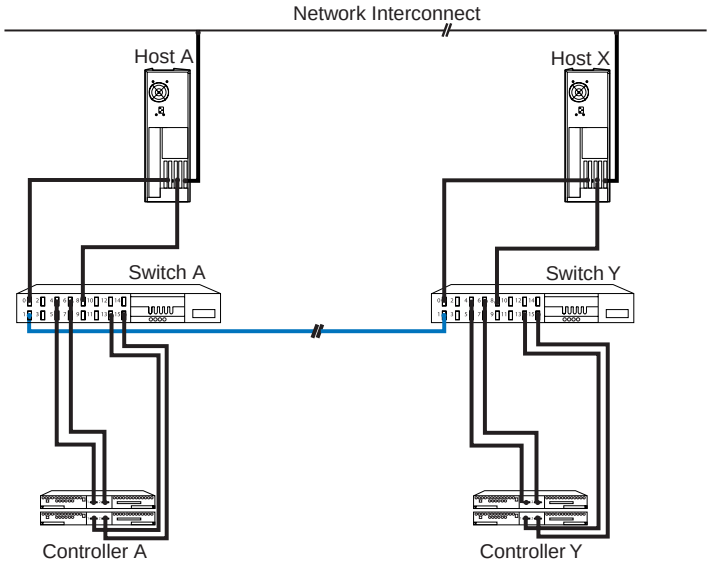

Figure 11: Single Fabric

Supported Hardware and Software

The hardware and software components required to implement data replication are segmented between the initiator (local) and target (remote) sites. Depending on the site configuration, each subsystem uses a minimum of one storage array. This section describes the supported hardware and software that DRM requires. For more detailed information, refer to the *Compaq Data Replication Manager Design Guide White Paper* available at

www.compaq.com/products/storageworks/Storage-Management-Software/DataRepindex.html

Supported Servers and Operating Systems

DRM supports the following servers with the appropriate operating systems and required software. For the most current list, go to the DRM home page at

www.compaq.com/products/storageworks/Storage-Management-Software/DataRepindex.html

Table 2. Supported Servers

Vendor	Operating Systems
Compaq–AlphaServer?TM	OpenVMS and Tru64
Compaq–ProLiant	Windows 2000, Windows NT, and NetWare
IBM–RS6000	AIX
Sun–Ultra SPARC and UE	Solaris

Table 3. Supported Operating Systems

Vendor	Operating System Versions
Compaq Open VMS	7.2-1H1, 7.2, 7.3
Compaq Tru64 UNIX	5.1, 5.1A
IBM AIX	4.3.3
Microsoft Windows NT	4.0 SP6a
Microsoft Windows 2000	Server, Advanced Server, Datacenter with SP2
Novell NetWare	5.1
SUN Solaris	7 (32- and 64-bit modes) 8 (32- and 64-bit modes)

Fibre Channel Switches

DRM supports the following Compaq Fibre Channel switches with all intersite link technologies listed in this white paper except where noted:

- Fibre Channel SAN Switch 8
- Fibre Channel SAN Switch 8-EL
- Fiber Channel SAN Switch 16
- Fiber Channel SAN Switch 16-EL
- StorageWorks SAN Switch Integrated/32 (not supported with DRM over ATM)
- StorageWorks SAN Switch Integrated/64 (not supported with DRM over ATM)

Storage Arrays

DRM supports the following Compaq Fibre Channel-based storage systems using the HSG80 controller. These storage arrays provide the investment protection needed for today's rapidly changing storage environments. Each HSG controller in the DRM solution must run ACS 8.6-1P controller firmware.

- EMA16000—Supports up to 168 disk drives and four pairs of controllers in one cabinet, yielding a storage density of more than 1 terabyte per square foot.
- EMA12000—Supports up to 126 disk drives and three pairs of controllers or 82 drives and one pair of controllers in one cabinet, depending on the configuration.
- MA8000—Supports up to 42 disk drives and one pair of controllers in an office cabinet.

- ESA12000—Supports up to 72 disk drives per pair of controllers.
- RA8000—Supports up to 24 disk drives per pair of controllers in an office cabinet.

Hard Drive

DRM supports the following disk drives when installed in the appropriate storage cabinet.

Table 4. Supported Compaq Universal Hard Drives

Capacity (gigabytes)	Size	Ultra 2 SCSI	Ultra3 SCSI	7,200 RPM	10,000 RPM	15,000 RPM
9	1 inch	Yes	Yes	Yes	Yes	Yes
18	1 inch	Yes	Yes	Yes	Yes	Yes
36	1 inch	N/A	Yes	N/A	Yes	Yes
72	1 inch	N/A	Yes	N/A	N/A	Yes
72	1.6 inch	N/A	Yes	N/A	Yes	N/A

N/A–Not available in that form.

Required Software

The following software is required to use the DRM solutions and is dependent on the operating system.

Table 5. Required Software

Operating System	Required Software
Compaq OpenVMS	MA/EMA OpenVMS platform kit DRM Solution kit
Compaq Tru64 UNIX	MA/EMA Tru64 UNIX platform kit DRM Solution kit
IBM AIX	MA/EMA Sun Solaris platform kit DRM Solution kit Secure Path for IBM AIX
Microsoft Windows 2000 and Windows NT	MA/EMA Windows NT platform kit DRM Solution kit Secure Path for Windows or HA/F500 Enhanced DT Cluster software
Novell NetWare	MA/EMA Sun Solaris platform kit DRM Solution kit Secure Path for Novell NetWare
Sun Solaris	MA/EMA Sun Solaris platform kit DRM Solution kit Secure Path for Sun Solaris

Glossary

Asynchronous Transfer Mode (ATM)—Communications networking technology for LANs and WANs that carries information in fixed-size cells of 53 bytes (5 protocol and 48 data).

Bi-directional DRM—A tested set of storage hardware and software products configured to allow two sites to use each other to maintain synchronized remote copies of online data.

Constant Bit Rate (CBR)—Category of ATM service that supports a constant or guaranteed data rate. CBR supports applications that need a highly predictable transmission rate.

Dual Switch/Single Site—A complete storage offering based on the DRM Dual Switch/Single Site for the Windows NT solution that provides the highest level of data availability in single-site environments. Based on switched fiber, this solution enables a customer to pinpoint and support a critical application and to scale up as that application grows or new applications are added.

Disaster Tolerance (DT)—Provides the ability for rapid recovery of user data from a remote location when a significant event or disaster occurs at the primary computing site.

Dense Wavelength Division Multiplexing (DWDM)—The ability to have multiple optical signals share a single optical cable.

E1—The standard European carrier for transmission at 2.048 Mb/sec.

E3—The standard European carrier for transmission at 34.304 Mb/sec.

Fabric—A network of Fibre Channel switches and attached devices.

Fibre Channel (FC)—Technology for very high speed, switching-based serial transmissions.

Gigabit Interface Converter (GBIC)—The hardware devices inserted into the ports of the Fibre Channel switch that hold the Fibre Channel cables. GBIC devices are available for short-range applications (0.5 to 500 meters), long-range applications (up to 10 km), and very long distances (up to 100 km).

Gigabit Link Module (GLM)—Permanently installed GLMs provide fiber-optic cable transmission at distances of 0.5 to 500 meters, depending on the cable size and quality of its installation.

Hop—One interswitch link.

ISL—Intersite link.

Initiator—For subsystems using the disaster-tolerant Data Replication Manager solution, the initiator is the storage array that is the primary source of information. In the event of a system outage, the database would be recovered from the target system.

Link—A connection between two adjacent Fibre Channel ports consisting of a transmit fiber and a receiving fiber. An example is the connection between the Fibre Channel switch port and the HSG80 controller.

MB/sec—Megabyte per second, sometimes shown as MBps.

Mb/sec—Megabit per second, sometimes shown as Mbps.

MMF—Multimode fiber, typically 50 micron, although 62.5 micron is also supported at reduced distances.

OC3—The optical carrier that provides high-speed bandwidth at 155.3 Mb/sec.

OSG—Open Systems Gateway. OSG converts Fibre Channel to and from ATM.

Peak Cell Rate (PCR)—Peak cell rate is the maximum transmission speed of a virtual connection and is a required parameter for the CBR service category.

Permanent Virtual Circuit (PVC)—Logical connection between two points that are manually defined by the network administrator.

Quality of Service (QoS)—Each virtual connection in an ATM network has a service category. The performance of the connection is measured by the established QoS parameter, which is outlined by the ATM forum.

Small Computer System Interface (SCSI)—An American National Standards Institute (ANSI) interface standard defining the physical and electrical parameters of a parallel I/O bus used to connect initiators to devices. A processor-independent standard protocol for system-level interfacing between a computer and intelligent devices including hard drives, disks, CD-ROMs, printers, scanners, and other devices.

Single-Mode Fiber (SMF)—A single-mode fiber, typically 9 micron, although 8 and 10 micron are also supported.

Synchronous Optical Network (SONET)—An ANSI standard for transmitting bits over fiber-optic cable.

Target—For subsystems using the disaster-tolerant Data Replication Manager solution, the target is the storage array that is the secondary or backup source of information. In the event of a system outage, the database would be recovered from the target system.

T1—The standard North American carrier for transmission at 1.544 Mb/sec.

T3—The standard North American carrier for transmission at 44.736 Mb/sec.

Unspecified Bit Rate (UBR)—Offers no traffic-related service guarantees.

Virtual Channel (VC)—The lowest-order logical address in ATM. VC refers to a given circuit on a link.

Virtual Channel Identifier (VCI)—The field of the cell header storing the VC address.

Virtual Path (VP)—The highest-order logical address in ATM. VP refers to a given group of circuits on a link.

Virtual Path Identifier (VPI)—The field of the cell header storing the VP address.

VLD GBIC—Very Long Distance GBICs.